

Claims

1. A method for communication between a network and a host computer having a processor and a sequential stack of protocol layers, the method comprising:

receiving, by the host from the network, a message packet including data and a plurality of header layers corresponding to the stack of protocol layers, said data intended for placement in the host according to protocol processing of said header layers,

processing said plurality of header layers with hardware circuitry,

including creating an indication of whether said packet is a candidate for a fast-path transfer of said data to a destination in said host, and

choosing, including referencing said indication, whether to process said packet by the protocol layers or to transfer said data according to said fast-path without processing said packet by the protocol layers.

2. The method of claim 1, wherein the host includes a first portion and a second portion that are connected by a PCI bus, with said hardware circuitry being disposed in said first portion and said protocol layers being disposed in said second portion, such that sending information between said first portion and said second portion includes transferring said information over said PCI bus.

3. The method of claim 2, further comprising creating, with said protocol layers, a communication control block for a connection including said packet, and transferring said communication control block from said second portion to said first portion, whereby said communication control block provides a guide for said fast-path.

4. The method of claim 3, further comprising transferring said communication control block from said first portion to said second portion, and processing said packet by said protocol layers.

5. The method of claim 1, further comprising creating a communication control block for a connection including said packet, wherein said communication control block provides a guide for said fast-path.

6. The method of claim 1, further comprising creating a communication control block for a connection including said packet, wherein said choosing includes comparing said indication with said communication control block.

7. The method of claim 1, further comprising creating a communication control block for a connection including said packet, wherein transferring said data to said destination according to said fast-path includes referencing said communication control block.

8. The method of claim 1, further comprising creating a communication control block for a connection including said packet, and transmitting, from the host to the network, a second message packet containing additional data and additional headers by referencing said communication control block.

9. The method of claim 1, further comprising:

processing said packet by the protocol layers, including sending said data to said destination,
receiving by said host from said network a second message packet containing additional data and additional headers, and transferring said second additional data to said destination, including referencing said communication control block.

10. The method of claim 1, wherein said destination is a file cache in the host.

11. The method of claim 1, wherein the host is connected to the network with a network interface device, and said processing said plurality of header layers is performed by said network interface device.

12. The method of claim 1, wherein said indication includes information regarding a transport layer header of said headers.

13. A method for communication of a host computer having a processor, a memory and a sequential stack of protocol layers, the host connected to the network with a network interface device, the method comprising:

receiving, by the network interface device from the network, a packet including data and a plurality of headers corresponding to the stack of protocol layers, said packet destined for the host,

categorizing said packet with the network interface device, including creating a status of said packet, and

choosing, including referencing said status, whether to process said packet with the stack of protocol layers or to bypass the stack of protocol layers and transfer said data to a destination in the host.

14. The method of claim 13, further comprising:

sending said packet to the stack of protocol layers,
processing said packet with the stack of protocol layers and thereby obtaining said destination for said message,
receiving, by the network interface device from the network, a related packet including additional data and additional headers, and
sending said additional data to said destination.

15. The method of claim 13, further comprising creating a context for a message including said packet, said context defining a connection between the host and a remote host, wherein choosing whether to process said packet with the stack of protocol layers or to bypass the stack of protocol layers includes comparing said status with said context.

16. The method of claim 13, further comprising creating a context for a message including said packet, said context defining a connection between the host and a remote host, and transferring said context from the host to the network interface device.

17. The method of claim 16, further comprising transferring said context from the network interface device to the host.

18. The method of claim 13, further comprising bypassing the stack of protocol layers by sending said data without said headers to said destination in a form suitable for said destination.

19. The method of claim 13, further comprising:
creating a context for a message including said packet, said context defining a connection between the host and a remote host, and employing said context for transmitting a reply to said network from said host, including prepending a header to reply data, said header including control information corresponding to said protocol layers.

20. The method of claim 13, wherein said destination is a file cache in said host.

21. The method of claim 13, wherein the network interface device includes logic circuitry adapted to analyze a transport layer header of said headers, and said categorizing includes analyzing said transport layer header with said logic circuitry.

20. The method of claim 13, wherein said status includes information regarding a transport layer header of said headers.

21. The method of claim 13, wherein said destination is a file cache.

22. A method for communication between a network and a host computer having a processor and a stack of protocol layers, the method comprising:

receiving, by the host from the network, a message packet including data, a plurality of headers corresponding to the stack of protocol layers and a checksum, wherein said data has been sent to the host for placement in the host according to protocol processing of said headers,

categorizing said packet to obtain a status of said packet, including analyzing said headers and checking said checksum, and

choosing whether to process said packet by the protocol layers or to send said data without said headers to a destination in said host without processing said headers by the protocol layers, said choosing dependent on said status.

23. The method of claim 22, further comprising;
creating a context for a message including said packet, and
transferring said data to said destination in accordance with said context.

24. The method of claim 22, wherein said analyzing includes processing a transport-layer header of said plurality of headers.

25. The method of claim 22, wherein the host is connected to the network with a network interface device, and said categorizing is performed by said network interface device.

26. The method of claim 25, further comprising creating a connection context for said packet with the protocol layers, and transferring said connection context to said network interface device.

27. The method of claim 25, further comprising transmitting, from said network interface device to the network, a second message packet, including creating a header for said packet by referencing said connection context.